REMARKS/ARGUMENTS

Reexamination and reconsideration of this application as amended is requested. By this amendment, Claims 1-22 have been amended. After this amendment, Claims 1-22 remain pending in this application.

Claim Objections

(1) The Examiner objected to Claims 7, 14 and 22 because of various informalities. Applicants have amended line 3 of Claims 7 and 14 and line 4 of Claim 22 to replace "block" with "blocks".

In view of the amendment to Claims 7, 14 and 22, Applicants believe that the objection to Claims 7, 14 and 22, as discussed above, has been overcome. Applicants request that the Examiner withdraw the objection to Claims 7, 14 and 22.

Claim Rejection under 35 U.S.C. § 112, second paragraph

(2) The Examiner rejected Claim 14 under 35 U.S.C. § 112, second paragraph, as being indefinite for failing to particularly point out and distinctly claim the subject matter which Applicants regard as the invention.

Applicants have amended claim 14 to more clearly and affirmatively recite "pinned memory buffers" and "a memory pool". This more clearly and affirmatively recites that "a system for allocating memory for use by time sensitive data communications processing" comprises, among other things, "pinned memory buffers" and "a memory pool". Support for this amendment may be found in the specification as originally filed. See for example page 6, lines 29 to 30, and page 10, line 29. No new

matter was added.

In view of the amendment to Claim 14 and the remarks above, Applicants believe that the rejection of Claim 14, under 35 U.S.C. § 112, second paragraph, as discussed above, has been overcome. Applicants request that the Examiner withdraw the rejection of Claim 14.

Claim Rejections - under 35 USC § 102

(3) The Examiner rejected Claims 1-6 and 8-13 and 15-21 under 35 U.S.C. 102(e) as being anticipated by Goldstein et al. (U.S. Patent 6,247,105).

Applicants have amended Claim 1 to more clearly and distinctly recite the present invention. The claim now more clearly and distinctly recites, "pinned kernel memory"; "a base pinned kernel memory block" and allocating "an additional pinned kernel memory block". Support for this amendment may be found in the specification as originally filed, see for example from page 6, line 19, to page 7, line 3, and page 10, lines 15-18. No new matter was added.

Goldstein et al.

Goldstein et al. discloses a method and apparatus for allocating memory in a computer system and for providing information relating to memory allocation. The invention disclosed in Goldstein et al. allows information relating to memory allocation for particular code to be obtained by external processes without having to modify that code. A memory space allocator maintains information in a data structure and updates the data structure with every allocation or deallocation of memory space. The area of memory that the memory space allocator has control over is referred to as an arena. The

memory area occupied by the arena is reserved by the operating system kernel for use only by the memory allocator.

The memory space arena contains a descriptor block and a number of memory blocks, which are comprised of memory pages. Each memory page is comprised of units of memory space, referred to as buckets, which vary in size. Preferably, the descriptor block is contained on a single memory page and comprises memory allocation information. The memory allocation information includes information as to the number of different types of memory blocks that have been defined within the memory space arena. Additionally, the bucket sizes within each memory block and the number of memory pages that have been instantiated within each memory block are also part of the memory allocation information.

When a memory page is instantiated, it is preferably mapped as a page in the virtual memory system and is capable of being swapped in or out of the system memory. When a page is swapped out of the system memory, it is stored on a hard drive. When the memory page is swapped into the system memory, it is stored in the system RAM. The process of swapping memory pages or the data contained therein is known in the art as paging.

When a process requests memory, the method disclosed by Goldstein et al. searches for a memory page that has the optimum number of free buckets of the desired size. A single bucket from the located memory page is then allocated to the process. The method then inquires to find out if sufficient memory space has been allocated to the process. If sufficient memory has been allocated the method ends, however, if more memory is needed the method continues. The method tries to allocate additional buckets from the same memory page that the first bucket was allocated from. However, if all of the buckets from the selected memory page have already been allocated, the method allocates across page boundaries (See for example, Col. 10, Lines 55-60).

Claim 1 Of The Present Invention

As now recited in amended Claim 1, the present invention, among other things, recites a method for allocating pinned kernel memory for use by time sensitive data communications processing. Implementing "pinned" kernel memory prevents the data stored therein from being paged out of real memory and into virtual memory (See for example, Specification, Page 3, Lines 10-27). The memory is global static fixed memory, namely such data structures as global variables reside in the pinned kernel memory and paging of the memory does not occur. Consequently, the data stored in the pinned kernel memory is available for immediate access by the processing hardware. Paging would not allow immediate access to the time sensitive information because the data would have to be swapped out of virtual memory first and into the system memory. Immediate access is crucial for the types of time sensitive information that the present invention is processing.

As recited for the method of amended Claim 1, a base pinned kernel memory block is established. The base pinned kernel memory block is allocated of an initial size and is not necessarily a maximum size to accommodate worst-case scenario memory requirements. More specifically, the size of the base pinned kernel memory block is not the total amount of pinned kernel memory that is allowed to be allocated to a process. As will be discussed below, the amount of pinned kernel memory that is available to a process is variable. A request for a pinned kernel memory buffer of a specific size is sent, for example, by a process and accepted by the method recited for Claim 1. Depending on the requested buffer size, sufficient pinned kernel memory may not exist within the currently allocated base pinned kernel memory block. If the method of amended Claim 1 determines that there is insufficient pinned kernel memory, the method allocates an additional block of pinned kernel memory from the unallocated available memory space. Therefore, the amount of pinned kernel memory that is available to a process is variable and not fixed. More specifically, the total amount of pinned kernel

memory available to a process is managed and not pre-allocated, it is allocated on an as needed basis.

The size of the additional pinned kernel memory block is at least the size of the requested buffer size. This ensures that the additional pinned kernel memory block will contain enough pinned kernel memory to satisfy the request. Also, the additional pinned kernel memory block is not required to form a contiguous pinned kernel memory section with the base pinned kernel memory block.

Contrast of Goldstein et al.

In contrast, Goldstein et al. does not disclose pinned kernel memory as now recited in amended Claim 1. Although the memory space in Goldstein et al. is reserved by the operating system, the operating system only reserves the memory space so that other processes do not interfere with it. Pinned kernel memory, on the other hand and as recited by amended Claim 1, is the memory used by the operating system and is configured so that the data contained therein is not paged out of real memory (See, for example, the instant specification, on page 3, lines 10-27). Goldstein et al. does not disclose or teach that the memory cannot be paged. Goldstein et al. explicitly discloses that memory pages within the memory blocks can be swapped in and out of the system memory (See for example Col. 5, Lines 44-52). Consequently, any data that is stored on those memory pages is swapped out into virtual memory. The memory is not global static fixed memory because the memory can be paged. Therefore, the memory disclosed in Goldstein et al. is not pinned kernel memory as now recited by amended Claim 1.

Goldstein et al. does not disclose establishing a base pinned kernel memory block and allocating an additional pinned kernel memory block as now recited by amended Claim 1. As already discussed in the previous paragraph, Goldstein et al. does not disclose pinned kernel memory. However, even assuming arguendo that Goldstein et al. did disclose pinned kernel memory, which Applicants already have shown it does not, a

base pinned kernel memory block and allocating an additional pinned kernel memory block as now recited by amended Claim 1 is clearly not disclosed, anticipated, or suggested by the Goldstein cited reference. Goldstein et al. discloses pre-allocating a worst-case scenario (maximum amount required) amount of memory. More specifically, Goldstein et al. discloses pre-allocating the total amount of memory that will be available for a process right from the start. The base pinned kernel memory block of the present invention can be added upon by an additional pinned kernel memory block. However, because the invention disclosed in Goldstein et al. pre-allocates all of the requested (optimum size of) memory for a process, the size of the pre-allocated memory is fixed and does not increase. (See for example, Col. 10, Lines 55-60). Also, the Examiner has made the observation that the base memory block in Goldstein et al. refers to "a memory page that is selected, which contains the optimum number of free buckets according to a memory request" (See for example, Office Action, Page 3). Note that this pre-allocated block of memory is then managed via virtual memory requests (paging). Therefore, it should be clear that Goldstein et al. does not disclose establishing a "base pinned kernel memory block" and allocating an additional pinned kernel memory block as now recited for amended Claim 1.

Furthermore, Goldstein et al. clearly does not disclose allocating an additional pinned kernel memory block. As discussed above, Goldstein et al does not disclose "pinned kernel memory" and the total amount of available memory for a process is pre-allocated right from the start. The size of pre-allocated memory is fixed and doe not increase. Memory management does not occur by increasing the allocated memory because all of the requested (optimum amount of) memory is already pre-allocated. Memory allocation requests from a process are all allocated from the pre-allocated memory area by utilizing virtual memory and paging. That is, if additional memory is needed by a process, Goldstein et al. performs paging operations that swap memory pages out of (or in to) the pre-allocated system memory by exchanging pages with the virtual memory (i.e., exchanged with the typically much slower disk drive storage memory) (See for example, Col. 10, Lines 55-60). The memory disclosed in Goldstein

et al. is specifically configured to be pageable, therefore, making the invention disclosed in Goldstein not suitable for the processing of time sensitive data (See for example, Col. 5, Lines 45-52). This is one of the deficiencies that the present invention addresses and overcomes. The present invention does not perform paging because additional pinned kernel memory blocks are allocated from the unallocated memory space as needed. System operations utilizing pinned kernel memory do not experience the significant delays of paging memory based operations. This is a significant advantage of the presently claimed invention.

Goldstein et al. also does not disclose that the "additional pinned kernel memory block is not required to form a contiguous pinned kernel memory section with the base pinned kernel memory block." As discussed above, the additional pinned kernel memory block is not an additional memory page within the base pinned kernel memory block, as the Examiner suggests on page 3 of the Office Action. The additional pinned kernel memory block is comprised of additional pinned kernel memory allocated from the available memory space.

Also, as discussed above, Goldstein et al. pre-allocates all of the available memory for a process up front and the allocated memory is fixed in size. Additional system memory cannot be added to the pre-allocated memory block. The memory that is being allocated is located all in one main memory area and not from separate pinned kernel memory blocks (that are not necessarily contiguous memory) as now recited for amended Claim 1. Therefore, Goldstein et al. clearly does not teach, anticipate, suggest, the "additional pinned kernel memory block is not required to form a contiguous pinned kernel memory section with the base pinned kernel memory block."

Therefore, in view of the amendment and remarks above, Applicants believe that since Goldstein et al. does not teach, anticipate, or suggest, the presently claimed "pinned kernel memory"; "a base pinned kernel memory block" or the allocation of an "additional pinned kernel memory block", the rejection of Claim 1 under 35 U.S.C. 102(e) has been

overcome. The Examiner should withdraw the rejection of this claim.

With respect to Claim 2, Applicants have amended Claim 2 to more clearly and distinctly recite the present invention. Claim 2 now conforms to Claim 8 and more clearly and distinctly recites a "kernel memory allocation module". Support for this amendment may be found in the specification as originally filed, see for example page 9, lines 11-15 and page 10, lines 17-30. No new matter was added.

Claims 2-6 depend from amended Claim 1 and, since dependent claims recite all of the limitations of the independent claim, it is believed that, therefore, these dependent claims also recite in allowable form. Accordingly, the above arguments and remarks regarding amended Claim 1 are also relevant to and support the allowability of Claims 2-6.

With respect to Claim 8, Applicants have amended Claim 8 to more clearly and distinctly recite the present invention. The claim now more clearly and distinctly recites, "a pinned kernel memory"; "a kernel memory allocation module"; "a base pinned kernel memory block" and allocating "an additional pinned kernel memory block". Support for this amendment may be found in the specification as originally filed, see for example from page 10, line 17, to page 13, line 26. No new matter was added.

The above arguments and remarks regarding Claim 1 and more specifically, "pinned kernel memory"; "base pinned kernel memory block" and the allocation of an "additional pinned kernel memory block" are likewise applicable here in support of the allowability of amended Claim 8. These applicable arguments have already been presented above and will not be repeated here.

Therefore, in view of the amendment and remarks above, Applicants believe that since Goldstein et al. does not teach, anticipate, or suggest, the presently claimed "pinned kernel memory"; "a kernel memory allocation module"; "a base pinned kernel memory

block"; or the allocation of an "additional pinned kernel memory block", the rejection of Claim 8 under 35 U.S.C. 102(e) has been overcome. The Examiner should withdraw the rejection of this claim.

With respect to Claim 9, Applicants have amended Claim 9 to more clearly and distinctly recite the present invention. Claim 9 now conforms to Claim 6 and more clearly and distinctly recites that "a second additional pinned kernel memory block" is allocated in "response to a determination that there is insufficient pinned kernel memory within the base pinned kernel memory block and the additional pinned kernel memory block". Support for this amendment may be found in the specification as originally filed, see for example page 10, lines 20-23. No new matter was added.

Claims 9-13 and 15 depend from amended Claim 8 and, since dependent claims recite all of the limitations of the independent claim, it is believed that, therefore, these claims also recite in allowable form. Accordingly, the above arguments and remarks regarding amended Claim 1 are also relevant to and support the allowability of Claims 9-13, and 15.

With respect to Claim 16, Applicants have amended Claim 16 to more clearly and distinctly recite the present invention. The claim now more clearly and distinctly recites, "pinned kernel memory"; "a base pinned kernel memory block" and allocating "an additional pinned kernel memory block". Support for this amendment may be found in the specification as originally filed, see for example from page 6, lines 29-30 and page 10, line 29. No new matter was added.

The above arguments and remarks regarding Claim 1 and more specifically, "pinned kernel memory"; "base pinned kernel memory block" and the allocation of an "additional pinned kernel memory block" are also applicable here in support for the allowability of amended Claim 16. These applicable arguments have already been presented above and will not be repeated here.

Therefore, in view of the amendment and remarks above, Applicants believe that since Goldstein et al. does not teach, anticipate, or suggest, the presently claimed "pinned kernel memory"; "a base pinned kernel memory block" or the allocation of an "additional pinned kernel memory block", the rejection of Claim 16 under 35 U.S.C. 102(e) has been overcome. The Examiner should withdraw the rejection of this claim.

With respect to Claim 17, Applicants have amended Claim 17 to more clearly and distinctly recite the present invention. Claim 17 now conforms to Claim 8 and more clearly and distinctly recites a "kernel memory allocation module". Support for this amendment may be found in the specification as originally filed, see for example page 9, lines 11-15 and page 10, lines 17-30. No new matter was added.

Claims 17-21 depend from amended Claim 16 and, since dependent claims recite all of the limitations of the independent claim, it is believed that, therefore, these claims also recite in allowable form. Note that the above arguments and remarks regarding amended Claim 1 are also applicable to and support the allowability of Claims 17-21. These applicable arguments have already been presented above and will not be repeated here.

Allowable Subject Matter

(4) The Examiner objected to Claims 7, 14 and 22, as discussed above in Section (1), but these claims would be allowable if rewritten or amended to overcome the objection(s). Additionally, the Examiner objected to Claim 14, as discussed above in Section (2), but this claim would be allowable if rewritten to overcome the rejection(s) under 35 U.S.C. 112, second paragrpah and to include all limitations of the base claim and any intervening claims.

Applicants have amended Claims 7, 14 and 22 to comply with Examiner's request. Claim 14 now recites in independent form including all limitations of the base claim and any intervening claims.

Therefore, Applicants submit that Claims 7, 14 and 22 are allowable, and request that the Examiner allow these claims to issue.

Conclusion

The foregoing is submitted as full and complete response to the Official Action mailed April 8, 2004, and it is submitted that Claims 1-22 are in condition for allowance. Reconsideration of the rejection is requested. Allowance of Claims 1-22 is earnestly solicited.

No amendment made was related to the statutory requirements of patentability unless expressly stated herein. No amendment made was for the purpose of narrowing the scope of any claim, unless Applicants have argued herein that such amendment was made to distinguish over a particular reference or combination of references.

Applicants acknowledge the continuing duty of candor and good faith to disclosure of information known to be material to the examination of this application. In accordance with 37 CFR §§ 1.56, all such information is dutifully made of record. The foreseeable equivalents of any territory surrendered by amendment are limited to the territory taught by the information of record. No other territory afforded by the doctrine of equivalents is knowingly surrendered and everything else is unforeseeable at the time of this amendment by the Applicants and the attorneys.

The present application, after entry of this amendment, comprises twenty-two (22) claims, including six (6) independent claims. Applicants have previously paid for twenty-two (22) claims including five (5) independent claims. Applicants, therefore,

believe that an additional fee of **$86** for claims amendment is currently due. The Commissioner is authorized to charge the claims amendment fee of **$86**, or if this fee amount is insufficient or incorrect, then the Commissioner is authorized to charge the appropriate fee amount, to Deposit Account 09-0463.

If the Examiner believes that there are any informalities that can be corrected by Examiner's amendment, or that in any way it would help expedite the prosecution of the patent application, a telephone call to the undersigned at (561) 989-9811 is respectfully solicited.

The Commissioner is hereby authorized to charge any fees that may be required or credit any overpayment to Deposit Account 50-1556.

In view of the preceding discussion, it is submitted that the claims are in condition for allowance. Reconsideration and re-examination is requested.

Respectfully submitted,

Date: 7/8/04

By: Jose Gutman
Jose Gutman
Reg. No. 35,171

Please send all correspondence concerning this patent application to:

Jose Gutman, Esq.
FLEIT, KAIN, GIBBONS, GUTMAN
BONGINI & BIANCO P.L.
551 N.W. 77th Street, Suite 111
Boca Raton, FL 33487
Tel (561) 989-9811
Fax (561) 989-9812